

ETSCC009 - Hadoop MapReduce for Mobile Clouds

Abstract

The new generations of mobile devices have high processing power and storage, but they lag behind in terms of software systems for big data storage and processing. Hadoop is a scalable platform that provides distributed storage and computational capabilities on clusters of commodity hardware. Building Hadoop on a mobile network enables the devices to run data intensive computing applications without direct knowledge of underlying distributed systems complexities. However, these applications have severe energy and reliability constraints (e.g., caused by unexpected device failures or topology changes in a dynamic network). As mobile devices are more susceptible to unauthorized access, when compared to traditional servers, security is also a concern for sensitive data. Hence, it is paramount to consider reliability, energy efficiency and security for such applications. The MDFS (Mobile Distributed File System) [1] addresses these issues for big data processing in mobile clouds. We have developed the Hadoop MapReduce framework over MDFS and have studied its performance by varying input workloads in a real heterogeneous mobile cluster. Our evaluation shows that the implementation addresses all constraints in processing large amounts of data in mobile clouds. Thus, our system is a viable solution to meet the growing demands of data processing in a mobile environment.

Index Terms—Mobile Computing, Hadoop MapReduce, Cloud Computing, Mobile Cloud, Energy-Efficient Computing, FaultTolerant Computing



Maruthi Plaza 91/6, TC Palya Main road,
Next to RK Apartments, Ramamoorthy Nagar,
Bangalore-560025.



9543218650



ieeeprojects@eminent.in